

**APPLICATION FOR
U.S. LETTERS PATENT**

by

Paul Dent

for

**MULTI-STAGE CDMA SYNCHRONIZATION
WITH PARALLEL EXECUTION**

**BURNS, DOANE, SWECKER & MATHIS, L.L.P.
Post Office Box 1404
Alexandria, Virginia 22313-1404
(703) 836-6620**

Attorney Docket No. 027575-196

MULTI-STAGE CDMA SYNCHRONIZATION WITH PARALLEL EXECUTION

BACKGROUND

5 The present invention relates to mobile phones or radio apparatus for receiving spread spectrum signals, such as code division multiple access (CDMA) signals, with which time synchronization is achieved.

 Radiocommunication systems involve the transmission of information over an air interface, for example by modulating a carrier frequency with that information.

10 Upon reception, a receiver attempts to accurately extract the information from the received signal by performing an appropriate demodulation technique. However, in order to demodulate a received signal, it is first necessary to synchronize timing between the transmitter and the receiver. Different levels of synchronization may be required depending upon the design of the radiocommunication system. For example,
15 in most systems clocking differences between the transmitter and the receiver create differences in timing at a bit level. Moreover, in some radiocommunication systems, information is transmitted in bursts, sometimes referred to as "frames", which represent chunks of information that are independently detected and decoded. In these types of systems it is also desirable to locate the beginning of a frame, so that
20 information relevant to a particular receiver is isolated and demodulated. Likewise, some systems (e.g., time division multiple access or TDMA systems) further subdivide frames into timeslots to create channels that are time multiplexed with one another. In these systems it is further desirable to synchronize the receiver to the beginning of each timeslot.

25 Some systems provide channelization using a spread spectrum technique known as code division multiple access (CDMA). In CDMA systems, the information data stream to be transmitted is first coded or spread using a unique spreading code and then combined with a long PN-sequence or a shorter scrambling-sequence. In the latter case, the scrambling-sequences are planned from cell to cell so that neighboring cells
30 use different scrambling-sequences or scrambling-masks. The information data stream

and the PN-sequence or the scrambling sequence can have the same or different bit rates. The multiplication of the information data stream with the unique spreading code and long PN-sequence results in an output stream of chips. Thus, in CDMA systems, it is also desirable to synchronize the receiver to the chip boundaries.

- 5 To further understand the synchronization tasks associated with signal processing in a CDMA radiocommunication system, consider the following example. FIG. 1 illustrates the use of base stations to transmit radio waves to mobile users (mobile stations) in a cellular system. In a CDMA system, base station 10 can transmit signals to mobile stations 14 and 15 as a single (composite) signal. The signal directed
- 10 to mobile station 14 is typically coded with a short code that is orthogonal or mostly orthogonal to a short code that is used to code the signal directed to mobile station 15. These signals are then spread with a second code that is sometimes referred to as a long code, associated with base station 10. The sum of the two coded and spread signals is then transmitted by base station 10.
- 15 When mobile station 14 receives the composite signal, mobile station 14 multiplies the spread signal with the long code and the short code to recreate the signal directed to mobile station 14 and the signal directed to mobile station 15 is suppressed as interference noise. Similarly, mobile station 15 multiplies the spread signal with the long code and the short code assigned to mobile station 15 to recreate the signal
- 20 directed to mobile station 15 and the signal directed to mobile station 14 is suppressed as interference noise. The receivers associated with mobile stations 14 and 15 must have acquired the various levels of synchronization to the received signal which were described above, in addition to learning or knowing the applicable long and short codes, in order to implement despreading, demodulation and decoding of the
- 25 information residing in that signal. Many different techniques have been developed in order to acquire synchronization at each of the various levels. For frame synchronization these techniques typically depend, in large degree, on the frame structure and the manner in which overhead or control information is conveyed to the mobile station. Overhead information is usually provided on one or more broadcast
- 30 control channels which are transmitted by base stations using known channels to which

-3-

mobile stations can quickly lock onto and receive the overhead information including, among other things, information used to acquire frame synchronization with that base station. Those skilled in the art will appreciate that many radiocommunication systems have unsynchronized base stations, i.e., base stations which do not share a common
5 timing reference signal. Accordingly, frame synchronization is a process which needs to be performed, for example, at start-up (i.e., when a mobile is powered on), as a mobile station moves from cell to cell and when measuring on channels associated with neighboring cells as part of cell reselection procedures (e.g., to confirm that a mobile station is listening to a "best" serving base station).

10 Of course, as with most signal processing tasks performed by the receiver, reducing the delay associated with synchronization is important in improving the receiver's performance. Many types of communication services, in particular speech communication, are relatively delay intolerant. Thus, system designers are continuously seeking for ways in which to reduce the amount of time that it takes to
15 perform any given signal processing task, including time synchronization.

SUMMARY

A radio receiver employing a synchronization method according to the present invention is able more rapidly to find and acquire synchronisation with a CDMA signal
20 by means of at least a two step process in which, at a first step, a number of candidate synchronisation frequencies or timings are identified, followed by confirming one of the candidates as a correct synchronisation state at a second or final step. According to the present invention, the confirmation step may be performed at the same time as the step of identifying further candidates by processing the same received signal samples in
25 different ways.

In an exemplary implementation, a first correlation means correlates shifts of a stream of received signal samples using a correlation length over which the received signal does not drift significantly in phase, amplitude or timing, thus allowing coherent correlation. Coherent correlations are not in general expected to reach a sufficient
30 signal-to-noise ratio to unambiguously identify with adequate certainty that correct

synchronisation has been achieved. Consequently a number of coherent correlations corresponding to like timing postulates are further accumulated non-coherently in a number of bins, each bin corresponding to a timing postulate. Non-coherent accumulation involves adding the magnitudes or square magnitudes of the coherent correlations, where the squared magnitude is equal to the sum of the squares of the real and imaginary parts of the coherent correlation value.

When non-coherent or magnitude accumulation must take place for a prolonged time period in order to identify a likely candidate timing, and the time period is so long that a drift of timing may occur that is of the order of plus or minus one timing bin width or more, the present invention may employ a drift compensation type of accumulation described in U.S. Patent Application Serial number 08/768,975 to Paul W. Dent, filed 18 Dec 1996, the disclosure of which is hereby incorporated by reference herein.

An exemplary CDMA system according to the present invention, using a 4 megachip per second modulation, searches time bins that are one chip wide, i.e. 0.25US. The receiver time and frequency reference has an initial error of +/- 10 parts per million, which results in a drift of one bin per 25mS. This exemplary CDMA system furthermore transmits a known chip pattern every 0.625mS for use by the receiver to achieve a first synchronization step. The known chip pattern has a length of 256 chips for example. With 10ppm frequency error, the maximum coherent correlation length is restricted to about 64 chips. Therefore four, 64-chip coherent correlations are accumulated non-coherently every 0.625mS to determine a correlation value with the known 256-chip pattern. Approximately 40 of these 0.625mS correlations may then be accumulated non-coherently, after which the timing may have drifted one chip. To compensate for such drift, after a number of non-coherent accumulations less than 40, for example 16, the results for each bin are accumulated with the best of previous cumulative results lying within +/-1 bin, thus allowing a drift of +/- one bin in 16 while still providing continued accumulation. In the above exemplary system, the number of timing bins is approximately $0.625\text{mS} \times 4\text{megachips/sec}$, that is about 2500 bins of one-chip width.

Drift-compensated accumulation continues according to the present invention until a cumulative results bin reaches a threshold. This indicates that the known signal pattern may have been found at the timing corresponding to the threshold. The timing associated with that bin is then recorded in a list of candidate timings to be further
5 evaluated, and the bin contents reset to zero. Drift compensated accumulation of correlations with the known signal pattern then continues and as further bins reach the first threshold, their associated timings are also entered into the list and their bin contents reset to zero. Eventually, a bin already reset to zero at least once may again reach the threshold value and would then be entered into the list a second time. The
10 list thus contains candidate timings in the order in which they reached the first correlation threshold, including possible repeats of an earlier candidate timing. Drift-compensated accumulation in timing bins can comprise subtracting the minimum value over all bins from all the bins, so as to emphasize differences between the bin values and prevent indefinite growth of numerical values. Detecting that a bin value reaches a
15 threshold can then imply detecting that a bin value has exceeded the other bin values by a threshold.

In the exemplary CDMA system, many signals may be transmitted in the same bandwidth at the same time using different spreading codes, preferably orthogonal codes. In addition to the above-mentioned known signal pattern, which typically uses
20 the same code, regardless of the transmitter, a second signal pattern is transmitted that is selected to be different for different transmitters in an adjacent group of transmitters. The second signal pattern being one of a limited number of, for example, 16 predetermined patterns. A second stage of synchronisation therefore consists of determining if one of the second known signal patterns can be found at a timing shift
25 recorded in the candidate list derived in the first step described above. Because the second known signal pattern is transmitted at the same time and on the same frequency as the first known signal pattern, they are both received overlapping at the receiver and correlations may be performed using the second known symbol pattern by a second correlation means simultaneously with continuing to perform and accumulate
30 correlations with the first known signal pattern using the first correlation means. The

invention of Dent and Wang described in U.S. Patent Application Serial Number 08/967,444, entitled "Efficient Correlation Over a Sliding Window" (filed November 11, 1997) may, for example, advantageously be employed to combine the first and second correlation means taking advantage of common computations to reduce effort and therefore power consumption.

Correlations are performed with all 16 of the second known signal patterns, but using only that timing shift identified with the first candidate in the list. According to a second aspect of the present invention, correlations with the first known signal pattern may also be used to estimate a frequency error, which is also stored in the list against candidate timings. When using a candidate timing to test for correlations with the 16 second known signal patterns, the frequency error estimate may be used to compensate the received signal for phase drift, allowing a longer coherent second correlation to be performed. When one of the 16 second correlations has reached a second threshold, the timing, frequency error estimate and the second known symbol pattern giving that correlation may be recorded in a second list of candidates to be tested in an optional third synchronization step. The present invention can optionally comprise performing simultaneous correlations with the second known symbol patterns using more than one of the timings from the candidate list in parallel. The results would be accumulated in $16N$ bins, where N is the number of simultaneously tested candidates. When any one of the $16N$ correlations reaches the second threshold, the associated second known code, timing and frequency error are transferred to the second list. The second list thus contains second correlations in the order in which their accumulations first reached the second threshold.

In the exemplary system, the second known symbol pattern identifies a group of third known symbol patterns, one of which should be found in a third correlation step. The present invention may be applied again to search simultaneously for a third correlation while continuing first and second correlation accumulations. The third correlation is performed using all known symbol patterns in the group, which once again may, for example, contain 16 members. When a third correlation is identified, the method has thus disambiguated which of 256 different signal waveforms are in use,

i.e. the CDMA spreading code used by the transmitter has been narrowed down to one candidate. Using that code, an attempt is then made to decode a broadcast control channel emitted by the transmitter, successful decoding being indicated, for example, by a Cyclic Redundancy Check Code (CRC). This final step confirms that a valid
5 signal has been identified and synchronisation achieved.

A further uncertainty that is resolved by the method is which of a number of frequency channels the transmitter is using. This is done by applying the method sequentially to all possible frequency channels selected intelligently in a priority order. For example, a frequency channel on which sync was previously found most recently
10 can be tested first. Moreover, sync search on any frequency channel may be abandoned at any stage if the correlation accumulation for that stage fails to reach the threshold for that stage in a given time.

BRIEF DESCRIPTION OF THE DRAWINGS

15 The features, objects, advantages and other aspects of the present invention will be more readily understood upon reading the following detailed description in conjunction with the figures, wherein:

Figure 1 illustrates an exemplary CDMA radiocommunication system in which the present invention may be employed;

20 Figures 2(a)-2(d) depict a flowchart which is used to described exemplary methods and systems for performing synchronization according to exemplary embodiments of the present invention; and

Figure 3 shows a CDMA signal format to which the present invention may be applied.

25

DETAILED DESCRIPTION

In the following description, for purposes of explanation and not limitation, specific details are set forth, such as particular circuits, circuit components, techniques, etc. in order to provide a thorough understanding of the present invention. However, it
30 will be apparent to one skilled in the art that the present invention may be practiced in

other embodiments that depart from these specific details. In other instances, detailed descriptions of well-known methods, devices, and circuits are omitted so as not to obscure the description of the present invention.

In Figure 2(a), a radio receiver, such as a mobile phone, performs a scan of the
5 allocated frequency band in step 1 upon first being switched on after a period of
inactivity. The frequency scan step identifies channels on which RF energy is detected
and can measure the signal strength of detected RF energy. In step 2, other
information is used to determine a prioritized list of channels that will be searched in
order. In the absence of other information, the prioritized order can be signal strength
10 order from strongest to weakest. Other information that can change this order would,
for example, include knowledge of channels on which valid signals had been found
before.

For example, if the signal strength now received on a channel on which sync
was established most recently is similar to that measured when valid synchronisation
15 was last achieved on that channel, it may be given highest priority. If the most recent
sync channel does not have the highest signal strength, but the signal strength on a
different channel on which sync is often established is high, then the different channel
may be given priority. Another example of other information which may affect the
scanning order is a list of channels on which sync has never been established or never
20 been established even when signal strength is high (because, for example, they might
be allocated to a different type of system). These channels may be given lower priority
than their measured signal strength would otherwise indicate. Another reason why
channels may be given low priority is when receiver frequency step sizes are narrower
than the system's channel spacing. For example, receiver resolution could be 200KHz
25 while a particular system places signals centered only on every 25th 200KHz step. A
radio receiver could learn on which of the 25 steps sync is most often achieved and
give those higher priority for searching than the other 24, depending of course on
whether the scan for energy confirms the likelihood that these steps are still those to
use. All such logic can be used in step 2 to establish a prioritized frequency channel
30 list.

In step 3, the receiver tunes to the next channel in the list, which is the first channel on the first pass through the flow chart. In step 4, the signal received on the tuned channel is processed using, for example, a correlator or matched filter matched to a known code. Exemplary matched filter structures are disclosed in U.S. Patent
5 Applications Serial No. 08/967,444, filed November 11, 1997 (Ericsson Docket #PO 8668) to Dent and Wang entitled "Efficient Correlation Over a Sliding Window" as described *supra*, which is hereby incorporated by reference herein. These inventive correlation techniques allow received signal samples to be correlated with a known code with all possible timing shifts between the known code and the signal samples
10 being correlated simultaneously.

Moreover, correlations can be performed using more than one code simultaneously with less than a commensurate increase in effort. Using these inventive correlators or matched filters, the savings in effort increases, the greater the number of correlations that have to be calculated. This property will prove useful in the current
15 invention which performs synchronisation in several stages using different codes, with all stages running in parallel.

The correlations computed in step 4 for different timing shifts between the received signal and the known code are accumulated with like values computed from a repetition of the same signal at periodic intervals. Such a signal format is shown in
20 Figure 3. A CDMA signal waveform allows several different signals to overlap in time and frequency by the use of different codes, which permit a receiver to separate the overlapping signals. In Figure 3, three types of signal are shown overlapping;

- (1) a Traffic Channel, which is divided into 10mS frame periods which are further subdivided into 16, 0.625mS timeslots;
- 25 (2) a Paging channel, which has the same frame and slot structure as the Traffic channel, although it may be of a different power level, and with other differences in the finer structure. One such difference is that the paging channel slots are divided into ten, 256-chip symbols, nine of which use a code unique to a particular transmitter selected from a set of

-10-

256 codes while the tenth symbol uses one of a subset of 16 codes and is called the second Periodic Channel or PERCH 2; and

- (3) a first Periodic Channel (PERCH 1) which comprises bursts of 256 chips using a code which is common to all transmitters. Moreover, the
5 PERCH 1 bursts are aligned in time with the PERCH 2 occurrences on the Paging Channel.

Optionally, one of the PERCH 1 or PERCH 2 occurrences may occasionally at regular intervals be an "unmodulated burst", such as would be achieved by setting all 256 chips to the same value, or systematically phase-rotating value, the latter giving a
10 frequency-offset CW burst. These CW bursts can be useful in providing the receiver with an opportunity to obtain coarse time and frequency synchronisation before beginning step 4 of Figure 2(a).

Techniques for utilizing such CW bursts in the context of TDMA signals (but not overlapping CDMA signals) are known from the European Digital Cellular System called GSM and, therefore are not elaborated here. The current description however
15 continues without reference to whether coarse synchronisation is achieved with the aid of CW bursts or not. Instead, the most onerous case is assumed, i.e., in which the frequency and time uncertainty over which the sync search is performed is not reduced. Consequently, in accordance with the exemplary waveform of Figure 3, correlation
20 with the PERCH 1 code can be found at any one of the 2560-chip shifts over the one-slot repetition period of the code.

It is also assumed that one 256-chip correlation will not yield an adequate YES/NO detection decision. An aim of the present invention is to permit a receiver to acquire time and frequency synchronisation even when the signal strength is too low to
25 establish a normal traffic channel, such as when a mobile phone user is deep inside a building. The aim of future cellular systems will be to provide a high-penetration method of alerting the phone that it is being called even in such disadvantaged circumstances, and conveying a short message, such as the caller-ID using highly redundant signaling. However, even using such highly redundant signaling,
30 synchronisation must first be established in order to decode the signal.

In order to improve the synchronisation robustness, corresponding correlations from corresponding timings in successive slots are accumulated in step 5. Since the timing uncertainty is 2560 chips in this example, 2560 timing bins are used to accumulate correlations corresponding to each particular timing postulate.

5 A correlation between a complex radio signal and a given code yields a complex number having an amplitude and a phase. The phase remains constant only over a limited time due to fading or frequency error. In an exemplary 2GHz radio system fading rates would allow correlations of length 256-chips or more to be made, but the result would exhibit a different phase angle from slot to slot due to the fading changing
10 the signal phase over the 0.625mS slot separation time. Consequently, only the magnitudes or magnitudes squared of the correlations are accumulated from slot-to-slot, which is termed "non-coherent" accumulation in the art.

 However, frequency error in the receiver due to the use of a low cost oscillator can restrict the coherency time to less than 256 chips. In an exemplary embodiment,
15 the coherent correlation length (matched filter length) is only 64 chips to allow for frequency error; four 64-chip correlations then have their magnitudes accumulated non-coherently to form a 256-chip correlation at step 4. Optionally, the four 64-chip correlations can be further processed to estimate a frequency error, as further elaborated later. If and when an opportunity arises to correct the oscillator's error, a
20 note can be made of that and subsequent 256-chip correlations can be coherent correlations.

 For example, the oscillator can be calibrated whenever synchronisation is confirmed with a valid transmitter and a note recorded against the temperature at which calibration was made. Subsequent sync searches made at temperatures for which the
25 oscillator has been calibrated can use 256-chip correlation while at other temperatures 64-chip coherent correlations would be made.

 At step 5, coherent or non-coherent correlations from step 4 are accumulated in the 2560 timing bins. If accumulation continues too long, frequency error in the receiver's chip rate relative to the transmitter's chip rate can cause a drift in the
30 correspondence of timing bins to a particular timing. This problem was solved by the

-12-

disclosure of U.S. Patent Application Serial No. 08/768,975 to Dent, filed December 18, 1996 which is hereby incorporated by reference herein. In the above application, drift is compensated by accumulating the largest of three adjacent previous cumulative values with a new correlation, thus allowing a timing drift of +/- one bin. In the
5 exemplary system, the frequency error is much smaller than that which would cause a drift of one chip per 2560 chips. A +/-10ppm frequency error will rather cause a 1 chip error in 100000 chips, or 40 slots of 2560 chips. Consequently, accumulation of corresponding values can take place in corresponding bins for up to 40 slots, and then the 2560, 40-slot, block-cumulative values are accumulated with the 2560 cumulative
10 bins using drift compensation according to the incorporated patent application. The need for 2 x 2560 memories can be avoided, by first processing the previous, cumulative bin values to determine the largest value in every adjacent group of three bins. The bin values are then replaced by this largest of three. Then each new correlation value for the next 40 slots is accumulated in a corresponding one of the
15 replaced bins. After 40 slots, the "best-of-3" operation is repeated prior to continuing. In practice, the exemplary system performs the "best-of-3" operation every 16 slots, corresponding to the exemplary 10mS frame period of Figure 3. In this way, cumulative correlation can continue indefinitely until something is detected, thus solving the "high-penetration sync" aim.

20 Even in the absence of a signal however, non-coherent accumulation of correlation magnitudes results in numerically growing values due to noise. To prevent indefinite numerical growth, it is useful to subtract the smallest value over all 256 bins from all the values, one bin then being zero and the non-zero bins indicating the margin by which they exceed the "noise" value. This differential value is in fact more useful
25 for detecting the presence of a signal. A predetermined threshold can be established against which the differential bin values are compared. If a bin value exceeds this threshold, it indicates that correlation with the associated timing value exceeded correlation with all other chip-shifts by a significant margin, indicating the likelihood, if not the certainty, that a signal exists with that timing.

One reason why certainty of detection does not exist with a single-stage sync detection algorithm is that the propagation path from a transmitter to a receiver can comprise many paths with different delays separated in time by one or more chips. Therefore each timing bin only accumulates that fraction of the energy received in a single ray; moreover multiple rays lead to multiple bins exhibiting correlations, so that there is uncertainty in which one to believe. The preferred algorithm is to continue drift-compensated accumulation of differential bin values until one bin value becomes the first to exceed the first detection threshold, as detected at step 6. If no bin exceeds the threshold after an accumulation cycle, at step 7 the receiver compares the time spent accumulating correlations on this frequency channel with a timeout limit. If the time limit is exceeded, further search is deemed useless and a return is made to step 3 to try tuning the next frequency channel in the list. The timeout limit can be adjusted in dependence on the measured signal strength values on the channel, so that a longer search time is allowed when the channel is weak. If all channels are weak, as measured at step 1, it can be an indication that the user is deep indoors and a much longer timeout used. Moreover, a different prioritization of the frequency channels to search can be employed if it is suspected that a "high-penetration" sync search must be performed.

If a bin exceeds the first detect threshold at step 6, step 8 of Figure 2(b) is entered. At step 8 of Figure 2(b), the just detected bin value is reset to zero and the timing (and optionally any frequency error that has been estimated during correlation) for that bin, i.e. the bin number, is recorded in a second list as a "candidate" or "possible" signal. Correlations then continue to be made and accumulated according to the method shown in Figure 2(a), except that changing frequency channel due to timeout is inhibited as long as the second list is not empty. Maintaining the same tuned frequency permits techniques according to the present invention to attempt a second stage of synchronisation on the channel on which first stage sync is detected. Although the inventive correlator of the above-incorporated applications can perform correlations of the same received signal samples with many different codes and code-shifts simultaneously, the receiver cannot receive signal samples on two or more different

frequency channels simultaneously, and so the choice of channel is frozen during performance of second stage of synchronization illustrated in Figure 2(b).

In Figure 2(b), the received signal is correlated at step 9 with each of 16 possible PERCH 2 256-chip codes. If coarse frequency error has been supplied from stage 1 sync detection, these 256-chip correlations may be coherent correlations. It will be explained later how the factor of four reduction in frequency error needed to increase the coherent correlation length from 64 to 256 chips is implemented.

The timing having been supplied from stage 1 sync, the timing uncertainty to be searched using the technique in Figure 2(b) is limited. For example, if the multipath delay spread is a maximum of 16 microseconds, only 64 timing bins are needed. However, there are 16 different codes, and 64 timing bins are needed for each, making 1024 bins altogether. At step 10, correlations are accumulated in bins corresponding to one of 16 codes and one of 64 timings for all candidates in the second list, thus there are multiple banks of 1024 bins required in the case of multiple candidates to be tested. Ultimately, the number of simultaneously processed candidates can be limited by the amount of memory. Candidates in the list that cannot be processed for lack of memory may be processed later after a previous candidate has been rejected at step 13.

Drift compensation may be applied as before using a "best of 3" operation every 1OmS frame. However, subtraction of the minimum value of these bins from all of them is not necessarily a good idea, as the timing has been so restricted that all timings and codes could conceivably contain valid signal energy arriving from more than one nearby base station transmitter. Therefore the minimum bin value from the bins of Figure 2(a) can be used as a noise floor and subtracted from all the bins of Figure 2(b) as well as all the bins of Figure 2(a). If any bin value in Figure 2(b) is then detected at step 11 to have exceeded this noise floor by more than a second threshold, step 14 is entered in figure 4. On the other hand, if no bin exceeds the second threshold, a test is made at step 12 to determine if enough time has been spent testing any of the timings in the second list. If so, the timing that has been tested for longer than the timeout period is deleted from the list. When all timings have been deleted from the list, the frequency channel change steps 7, 8 of Figure 2(a) may be re-enabled.

Referring now to Figure 2(c), step 14 is entered with a timing value and code giving a second threshold detection at step 11. This is recorded in a third list, allowing again for multiple candidates. The PERCH 2 code detected at step 11 indicates one of 16 groups of codes, each group containing 16 codes, that the transmitter may be using for coding the other nine symbols per slot of the paging channel. Consequently correlations are now performed using the timing and coarse frequency error from previous steps with the nine other paging channel symbols using each of the 16 codes in the indicated group. These are accumulated as before in a limited number of timing bins times the 16 code bins. When one of these third banks of bins is detected at step 17 to exceed a third threshold above the noise floor, the code used by the transmitter is uniquely suggested and is conveyed to a fourth list at step 20 of Figure 2(d) for confirmation. It may be seen that Figure 2(c) operates similarly to Figure 2(b) and so will not be elaborated further. A person skilled in the art will realize that the operations of flow charts of Figures 2(a)-2(d) can be controlled by a microprocessor (not shown) using a suitable software program, and that software programs that are identical apart from the data memory elements they operate on do not need to be stored in multiple copies, but can be given the property of re-entrancy that allows a single program to be multiply active in different states and quasi-simultaneously to process different data and even in different ways controlled by different parameters.

Now referring to Figure 2(d), step 20 is entered having identified one of in total 256 codes that the detected transmitter is thought to be using to code the paging channel signal. It may be that more than one candidate code has been detected, for example due to the mobile station being within range of more than one base station transmitter using the same frequency channel. Therefore the just-identified signal parameters are entered into a fourth list with any other codes, timings and frequency errors already in the list. At step 21 an attempt is made to decode the signal at the top of the list. If the signal is a CDMA signal, a RAKE receiver (not shown) may be appropriate to perform this decoding step. The RAKE receiver first performs correlations with known codes based on the identified unique code of this transmitter to establish the phase and amplitudes (i.e. complex channel parameters) of each significant

5 multipath ray reaching the receiver. The RAKE receiver then uses the code in correlation with unknown data symbols using the timing shifts of the established significant rays. The results of correlation with unknown data symbols are combined using weighting factors based on the established complex channel parameters for each ray. The combined value is then submitted as a "soft decision" to an error correction decoding algorithm to decode information transmitted on the paging channel. More details of RAKE receivers may be found in U.S. Patents Nos. 5,305,349 and 5,572,552 to Applicant which are hereby incorporated by reference herein.

10 After error correction decoding of the signal at step 21, a test for correct decoding is made at step 22. This can be done for example with the aid of a Cyclic Redundancy Check (CRC) code. If the CRC code does not check, this may be an indication that the signal in the fourth list was a false signal or of low signal quality. Rejection of the signal and deletion from the fourth list may occur as a result of several successive failures of the CRC check. As with all the steps of Figure 2(a)-2(d) presented here, attempted decoding of one, any or all of the signals with parameters recorded in the fourth list can in principle be done simultaneously, and while the earlier synchronisation detection stages are also still proceeding. In this context "simultaneously" means processing the same received signal in many different ways corresponding to searching for a PERCH 1 burst, a PERCH 2 burst or CRC check using different despreading codes. The term "simultaneously" does not necessarily imply multiple copies of decoding hardware, as sufficiently fast hardware can be time-shared, i.e. operated sequentially, to process the same stored data in many different ways. The amount of parallelism provided in the hardware to obtain the desired processing ability is a design trade-off that is not material to the inventive principles disclosed above.

25 Using the invention, a multi-stage synchronisation process is implemented with the aim of reducing the time needed to acquire synchronisation with a signal such as a spread-spectrum signal. The first stage of the sync process is designed to speed detection of any signal, by limiting the number of codes used for the first detection stage and accepting a higher probability of false alarms. Later stages of

synchronisation are performed using only the limited number of candidate signals defined by the parameters identified in the earlier stage, and aim to successively limit the number of signal candidates retained for further testing while improving the confidence that a retained signal is a true signal. The final step of successfully
5 decoding a paging channel broadcast confirms detection of a true signal. When a true signal is detected and decoded, further information may be read by the receiver regarding surrounding stations, frequencies and other parameters of the network. The receiver may also measure its frequency error relative to a true signal and store a frequency calibration value for the prevailing temperature, in order to learn a
10 frequency/temperature correction table for its reference oscillator with which to accelerate synchronisation on future occasions.

When initially a receiver has no preexisting frequency/temperature calibration table, coherent correlation lengths may be restricted, as described above. Coherent correlations of limited length have to be combined or accumulated non-coherently to
15 produce correlations over longer periods. This can be accomplished using non-coherent accumulation or combination devices to compute the magnitude (or magnitude squared) of a complex correlation value and combine the magnitudes. Non-coherent correlation will always produce a positive result that grows with time, and is therefore more liable to falsely exceed a detect threshold than a coherent correlation.
20 Consequently it is of interest to estimate frequency error from a preceding, non-coherent stage of correlation (e.g. that of Figure 1(a)) in order to permit a subsequent stage of correlation to be performed coherently by compensating for the estimated frequency error (e.g. the stage described by Figure 2(b)).

In the exemplary system, the synchronization codes transmitted are of length
25 256 while the maximum coherent correlation length imposed by frequency error is 64. To permit the full coherent correlation length of 256 to be used, it is therefore necessary to reduce the frequency error by a factor of four. This may be done in the following exemplary manner.

In Figure 2a), 64-point correlations are computed for four, successive 64-
30 symbol blocks over a 256-symbol block. Denoting these by the complex numbers

-18-

$$I1 + jQ1$$

$$I2 + jQ2$$

$$I3 + jQ3$$

and $I4 + jQ4,$

- 5 non-coherent combination of these would normally comprise computing
 $\text{SQRT}(I1^2 + Q1^2) + \text{SQRT}(I2^2 + Q2^2) + \text{SQRT}(I3^2 + Q3^2) + \text{SQRT}(I4^2 + Q4^2)$
 or alternatively the sum-of-squares
 $I1^2 + Q1^2 + I2^2 + Q2^2 + I3^2 + Q3^2 + I4^2 + Q4^2$

- Coherent combination, if the frequency error was zero, would comprise
 10 computing

$$(I1 + I2 + I3 + I4) + j(Q1 + Q2 + Q3 + Q4).$$

Consider now a frequency error that gives rise to a 90 degree phase rotation over each 64-symbol correlation interval, equivalent to a multiplication by "j" for a positive frequency error or by -j for a negative frequency error.

- 15 Then, for positive frequency error the correlations would correspond to

$$I1 + jQ1$$

$$jI2 - Q2 \quad (\text{having been multiplied by } j)$$

$$-I3 - jQ3 \quad (\text{having been multiplied by } j^2)$$

and $-jI4 + Q4 \quad (\text{having been multiplied by } j^3)$

- 20 The sum of the real parts is then $I1 - Q2 - I3 + Q4$
 and the sum of imaginary parts is $Q1 + I2 - Q3 - I4.$

For negative frequency error, the corresponding combinations would be

$$I1 + Q2 - I3 - Q4 \quad \text{and} \quad Q1 - I2 - Q3 - Q4$$

- For a frequency error, positive or negative, that gives 180 degrees phase shift
 25 per 64-symbol correlation interval, the corresponding accumulations would be

$$I1 - I2 + I3 - I4 \quad \text{and} \quad Q1 - Q2 + Q3 - Q4$$

Collecting all of the above, we obtain four different ways to accumulate four 64-symbol correlations:

- 1) $I1 + I2 + I3 + I4 \quad Q1 + Q2 + Q3 + Q4$ for small frequency error
 30 2) $I1 - Q2 - I3 + Q4 \quad Q1 + I2 - Q3 - I4$ for +90 degrees per 64 chips

-19-

- 3) $I2 + Q2 - I3 - Q4 \quad Q1 - I2 - Q3 + I4$ for -90 degrees per 64 chips
- 4) $I1 - I2 + I3 - I4 \quad Q1 - Q2 + Q3 - Q4$ for +/- 180 degrees per 64 chips

A person skilled in the art will recognize the above as computing a four-point FFT. If desired, the correlation results can be combined instead with phase rotations of
 5 +/-45 degrees and +/-135 degrees instead of 0, +/-90 and +/-180 degrees to avoid the frequency error sign ambiguity of the +/-180 degree combination. This would slightly modify the transition rules for drift compensated accumulation disclosed below, but nevertheless in a manner that may be fully derived by a person of average skill in the art with the help of these teachings.

10 An implementation of the current invention which provides the highest performance therefore comprises accumulating four 64-symbol correlations using all four of the above formulae. Then the magnitudes of the resulting 256-chip correlations are accumulated from one signal repetition period to the next (i.e. slot-to-slot in the exemplary format of Figure 3) in a set of 4 x 2560 bins, each set of 2560 bins
 15 corresponding to one of the above frequency error assumptions.

When a bin value exceeds a detect threshold, the bin number then yields both a timing value (1-2560) and a frequency error (1-4) to use for accumulating correlations in Figure 2(b). Note that it does not matter at this stage that with frequency error (4) we cannot tell if it is positive or negative. If the accumulation of type (4) above gave
 20 largest correlation for sync detect stage 1, then it should also give correct accumulation for sync detect stage 2.

It is seen that the above provides a means to classify frequency error into one of four smaller errors of $\frac{1}{4}$ the error, as needed to allow 256-symbol coherent correlations.

25 It is also possible to improve drift-compensated accumulation when frequency error is detected in the above way. As described above, a bin value may be periodically replaced by the largest of three adjacent values to allow for a one-bin timing drift over the period. If a bin is replaced by the next highest numbered bin value, it is an indication that the received frequency was

-20-

low relative to the receiver's frequency reference crystal oscillator and had drifted "late". Conversely, replacing a bin value by the next lower bin number indicates that the received frequency was high.

In the case of simultaneous frequency error detection using 4×256 bins, let the
 5 bins be numbered using two indices j, k with k ranging from 1 to 2560 and j ranging from -1 to +3 to indicate frequency nominal ($j=0$); frequency high ($j=1$); frequency low ($j=-1$) and frequency very high or low ($j=+3$).

Now, when a bin with second index k is replaced with a value from a bin with second index $k-1$, it indicates that the receiver's frequency was low compared to the
 10 received signal. Thus, to be consistent, this should only occur for a second index j equal to -1 or +3; conversely, replacing bin value index k with $k+1$ should only be allowed for $j = +1$ or +3. Replacing bin value index k with itself implies no frequency error, so should be done only for $j=0$.

In practice, due to frequency error not falling exactly into one of the four
 15 categories, some opportunity for "cross-fertilization" should be allowed, resulting in the following strategy:

- 1) Replace $\text{bin}(0,k)$ by the largest of
 $\text{bin}(0,k)$, $\text{bin}(-1,k-1)$ and $\text{bin}(1,k+1)$
 (frequency nominal, a bit low or a bit high)
- 20 2) Replace $\text{bin}(1,k)$ by the largest of
 $\text{bin}(1,k)$, $\text{bin}(1,k+1)$ and $\text{bin}(3,k+2)$
 (frequency more than a bit high)
- 3) Replace $\text{bin}(-1,k)$ by the largest of
 $\text{bin}(-1,k)$, $\text{bin}(-1,k-1)$ and $\text{bin}(3,k-2)$
 25 (frequency more than a bit low)
- 4) Replace $\text{bin}(3,k)$ by the largest of
 $\text{bin}(1,k+1)$, $\text{bin}(-1,k-1)$, $\text{bin}(3,k-2)$ and $\text{bin}(3,k+2)$
 (frequency very high or very low)

A person skilled in the art may devise other transition rules based on the same
 30 logic, namely that if symbol rates and radio frequency are derived at the transmitter

and receiver from the same reference oscillators, then the sign of timing drift must be related to the sign of frequency error. The above transition rules thus involve less computation than finding the largest of twelve values (4 frequency errors times 3 timing drifts of early, late and nominal), but may be restricted to finding the largest of
5 3 or 4 values.

The use of a four-point FFT for simultaneously accumulating correlations in four different ways corresponding to different frequency error postulates may of course be extended to larger FFTs for finer resolution of frequency error. An advantage of four-point FFTs however, which are sufficient for the preferred system parameters, is
10 that they do not require any multiplications but require only additions, which reduces the complexity of the inventive receiver according to the invention.

The above-described exemplary embodiments are intended to be illustrative in all respects, rather than restrictive, of the present invention. Thus the present invention is capable of many variations in detailed implementation that can be derived from the
15 description contained herein by a person skilled in the art. All such variations and modifications are considered to be within the scope and spirit of the present invention as defined by the following claims.